Generative AI Meets Computer Graphics

Lingjie Liu
University of Pennsylvania
Philadelphia, USA
lingjie.liu@seas.upenn.edu

Minghao Guo
Massachusetts Institute of Technology (MIT)
Cambridge, USA
guomh2014@gmail.com

Yin Yang
University of Utah
Salt Lake City, USA
yangzzzy@gmail.com

Chenfanfu Jiang
University of California Los Angeles
LA, USA
chenfanfu.jiang@gmail.com

Wojciech Matusik
Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT)
Boston, USA
wojciech@mit.edu

Rana Hanocka
University of Chicago
Chicago, USA
ranahanocka@uchicago.edu

Jiajun Wu
Stanford University
Stanford, USA
jiajunwu@cs.stanford.edu

Xin Tong
Anuttacon
Bay Area, USA
xtong.gfx@gmail.com

Yuanming Hu
Meshy AI
Bay area, USA
ethan@meshy.ai

Yifei Li
Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT)
Boston, USA
liyifei@mit.edu

Yao Feng
Stanford University
Stanford, USA
yaofeng@stanford.edu

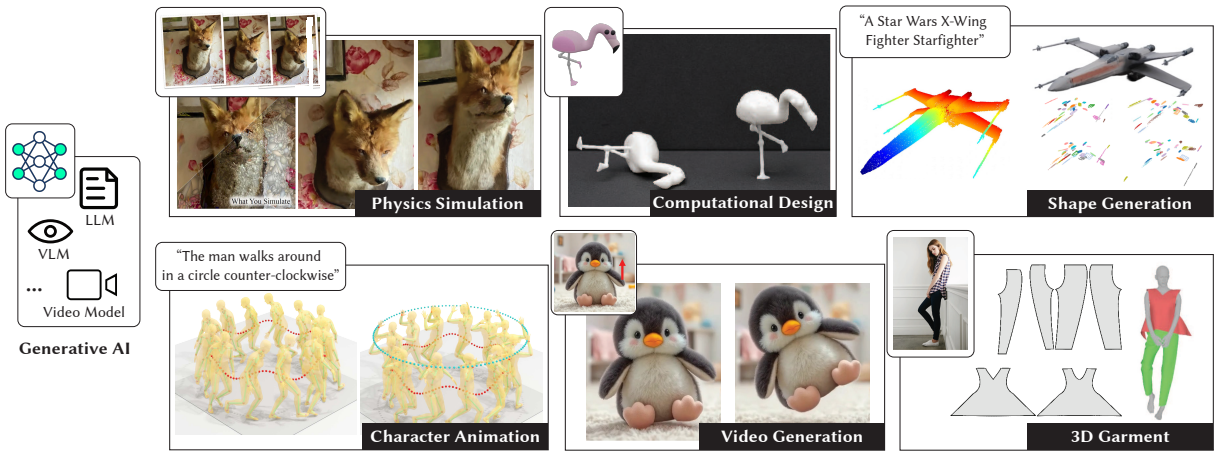

Figure 1: Generative AI is rapidly becoming a unifying interface for computer graphics, connecting LLM/VLM/video foundation models with geometry- and physics-aware representations. Our workshop highlights emerging directions spanning physics simulation, computational design, 3D shape generation, character animation, video generation, and garment synthesis, emphasizing multimodal control, editability, and physically plausible outputs. Example results are adapted from selected papers by the workshop speakers and organizers [Guo et al. 2024; Li et al. 2025; Wan et al. 2024; Wang et al. 2025; Xie et al. 2024; Yang et al. 2025].

Abstract

Generative AI meets Computer Graphics explores how modern generative models are reshaping the classical computer graphics pipeline, spanning 3D content creation, physics simulation, geometry processing, computational design, and animation. The workshop will bring together researchers and practitioners working on text-/image-conditioned 3D generation, learning-based and hybrid simulation methods, and generative tools for geometric editing, retopology, and asset refinement, as well as generative approaches for design and physically plausible motion and dynamics generation. Beyond showcasing recent advances, the workshop will emphasize practical integration into graphics production pipelines, discuss

core challenges such as controllability, editability, physical consistency, and evaluation, and highlight future directions that connect generative AI with the foundational problems of computer graphics more closely.

Keywords

generative AI, foundation models

ACM Reference Format:
Lingjie Liu, Minghao Guo, Yin Yang, Chenfanfu Jiang, Wojciech Matusik, Rana Hanocka, Jiajun Wu, Xin Tong, Yuanming Hu, Yifei Li, and Yao Feng. 2026. Generative AI Meets Computer Graphics. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Technical Workshops (SIGGRAPH Technical Workshops '26), July 19–23, 2026, Los Angeles, CA, USA*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3799828.3811423

1 Motivation

Generative AI is rapidly becoming a foundational tool for computer graphics. Recent advances in generative modeling have moved beyond text and 2D image synthesis to enable text-/image-conditioned 3D asset creation, learning-assisted physics simulation, and generative workflows for geometry processing, design, and animation. This progress is opening new possibilities for accelerating content creation and enabling more intuitive, high-level control over complex graphical processes.

At the same time, deploying generative methods in real graphics pipelines raises fundamental challenges that the community is actively grappling with: controllability and editability of generated content, physical plausibility and consistency over time, robust evaluation and benchmarking for 3D and dynamic outputs, and practical concerns around data quality, interoperability, and integration with existing production tools. The goal of this workshop is to provide a focused venue for researchers and practitioners to share recent technical advances, compare perspectives across subfields, and identify the most important open problems and future research directions. By connecting generative AI with the core foundations of computer graphics, such as representation, geometry, physics, and design, the workshop aims to help shape the next generation of methods and systems for creating and simulating digital worlds.

2 Workshop Structure

Opening and framing (10 minutes). The workshop will begin with a concise overview of the goals and scope of "Generative AI Meets Computer Graphics," establishing the motivation for bringing together perspectives from graphics, vision, and machine learning. This opening segment will introduce the key themes of the workshop, including generative 3D content creation, learning-based and hybrid simulation, controllability, editability, evaluation, and integration into practical graphics workflows.

Invited talks (5 talks, 30 minutes each). The main technical content will consist of five invited talks, each 30 minutes, representing complementary perspectives across the workshop themes. Together, these talks will cover recent research advances while emphasizing the representations, modeling choices, and system considerations needed to translate generative methods into practical graphics workflows, and will connect core ideas to constraints encountered in real production pipelines.

Lightning talks (2 talks, 15 minutes each). To broaden participation and highlight emerging directions, the workshop will include two lightning talks of 15 minutes each. This segment is designed for concise presentations of focused ideas, early-stage results, evaluation efforts, or pipeline-oriented systems work, helping surface new viewpoints that complement the invited program and broaden the discussion across the workshop themes.

Closing remarks (5 minutes). The workshop will conclude with brief closing remarks that synthesize recurring themes across the invited and lightning talks and highlight future directions for generative AI in computer graphics, including evaluation, controllability, physical plausibility, and integration into production workflows.

3 Schedule Overview

Time	Session
0:00 – 0:10	Opening
0:10 – 0:40	Invited Talk 1
0:40 – 1:10	Invited Talk 2
1:10 – 1:40	Invited Talk 3
1:40 – 2:10	Invited Talk 4
2:10 – 2:40	Invited Talk 5
2:40 – 2:55	Lightning Talk 1
2:55 – 3:10	Lightning Talk 2
3:10 – 3:15	Closing remarks

Total Duration: 3.25 hours.
All times approximate; subject to adjustment based on final speaker availability.

4 Target Audience

This workshop targets researchers, practitioners, and students in the computer graphics community who are interested in how generative AI is reshaping core graphics problems and workflows, including 3D content generation, physics simulation, geometry processing, computational design, and animation. It will be especially relevant to attendees developing generative models and hybrid learning-and-graphics methods, as well as those focused on making generative approaches controllable, editable, physically plausible, and compatible with established graphics pipelines.

The workshop also welcomes participants from adjacent SIGGRAPH communities, such as machine learning, computer vision, robotics, and HCI/interactive tools, who are working on representations of 3D geometry and dynamics, interactive authoring, and evaluation of 3D and dynamic generative outputs. Finally, the workshop is intended to attract industry researchers and engineers from film, games, AR/VR, and design/CAD, where generative AI is increasingly influencing production workflows and where practical constraints and user needs can help ground discussion and shape impactful future directions.

5 Speaker Biographies

Keynote Speakers:

- **Prof. Wojciech Matusik (MIT)** (senior, academic, confirmed) Wojciech is a Professor of EECS at MIT CSAIL, where he leads the

Computational Design and Fabrication Group and is a member of the Computer Graphics Group. Before coming to MIT, he worked at Mitsubishi Electric Research Laboratories, Adobe Systems, and Disney Research Zurich. He studied computer graphics at MIT and received his PhD in 2003. He also received a BS in EECS from the University of California at Berkeley in 1997 and MS in EECS from MIT in 2001. His research interests are in computer graphics, computational design and fabrication, computer vision, robotics, and HCI. In 2004, he was named one of the world's top 100 young innovators by MIT's Technology Review Magazine. In 2009, he received the Significant New Researcher Award from ACM SIGGRAPH. In 2012, he received the DARPA Young Faculty Award and he was named a Sloan Research Fellow.

- **Prof. Rana Hanocka (University of Chicago)** (junior, academic, confirmed)
  Rana is an Assistant Professor of Computer Science at the University of Chicago, where she founded and directs 3DL (threedle!), a group of enthusiastic researchers passionate about 3D, machine learning, and visual computing. She obtained her Ph.D. in 2021 from Tel Aviv University under the supervision of Daniel Cohen-Or and Raja Giryes. Her research is focused on building artificial intelligence for 3D data, spanning the fields of computer graphics, machine learning, and computer vision. While deep learning, the most popular form of artificial intelligence, has unlocked remarkable success on structured data (such as text, images, and video), she is interested in harnessing the potential of these techniques to enable effective operation on unstructured 3D geometric data. The 3DL group has developed a convolutional neural network designed specifically for meshes and has also explored how to learn from the internal data within a single shape (for surface reconstruction, geometric texture synthesis, and point cloud consolidation).
- **Prof. Jiajun Wu (Stanford University)** (junior, academic, confirmed)
  Jiajun is an Assistant Professor of Computer Science at Stanford University, working on computer vision, machine learning, and computational cognitive science. Before joining Stanford, he was a Visiting Faculty Researcher at Google Research. He received his PhD in Electrical Engineering and Computer Science at the Massachusetts Institute of Technology. His research has been recognized through the ACM Doctoral Dissertation Award Honorable Mention, the AAAI/ACM SIGAI Doctoral Dissertation Award, the MIT George M. Sprowls PhD Thesis Award in Artificial Intelligence and Decision-Making, the 2020 Samsung AI Researcher of the Year, the IROS Best Paper Award on Cognitive Robotics, and faculty research awards and graduate fellowships from JPMC, Samsung, Amazon, Meta, Nvidia, and Adobe.
- **Dr. Xin Tong (Anuttacon)** (senior, industry, confirmed)
  Xin is a Principal Researcher at Anuttacon. Prior to Anuttacon, he spent 25 years at Microsoft Research Asia (1999–2024), where he worked on computer graphics and computer vision, including 3D digitization and related visual computing problems. He received his PhD from Tsinghua University and his MS/BS from Zhejiang University. His contributions have been recognized by the 2024 Asiagraphics Outstanding Technical Contributions Award.
- **Dr. Yuanming (Ethan) Hu (Meshy AI)** (junior, industry, confirmed)
  Yuanming (Ethan) is the Co-founder and CEO of Meshy AI, a generative AI product focused on 3D creation. Before founding Meshy, he earned his Ph.D. in MIT EECS at MIT CSAIL, advised by Frédo Durand and Bill Freeman. His work spans computer graphics, AI, and high-performance systems; he is also known as the original author of the Taichi programming language. His PhD thesis received an Honorable Mention in the SIGGRAPH 2022 Outstanding Doctoral Dissertation Award.

Lightning Talk Speakers:

- **Yifei Li (MIT)** (junior, academic, confirmed)
  Yifei is a final-year PhD student at MIT CSAIL in the Computational Design and Fabrication Group, advised by Prof. Wojciech Matusik. She received her bachelor's degree in Computer Science from Carnegie Mellon University. Her research builds intelligent physical systems that jointly optimize form and control, and learn from data through physics-based models. She was selected as a Rising Stars in EECS, Rising Stars in Computer Graphics, and received the MIT Stata Family Presidential Fellowship.
- **Dr. Yao Feng (Stanford University)** (junior, academic, confirmed)
  Yao is a Postdoctoral Scholar in Computer Science at Stanford University. Her research focuses on capturing and understanding digital humans at scale, spanning computer vision, computer graphics, biomechanics, and robotics. She received her PhD from ETH Zürich and the Max Planck Institute for Intelligent Systems, under the supervision of Michael J. Black and Marc Pollefeys, and she is currently working at Stanford with Karen Liu (postdoctoral faculty sponsor), as well as Jennifer Hicks and Scott L. Delp.

References

Minghao Guo, Bohan Wang, Pingchuan Ma, Tianyuan Zhang, Crystal Owens, Chuang Gan, Josh Tenenbaum, Kaiming He, and Wojciech Matusik. 2024. Physically compatible 3d object modeling from a single image. *Advances in Neural Information Processing Systems* 37 (2024), 119260–119282.

Xuan Li, Chang Yu, Wenxin Du, Ying Jiang, Tianyi Xie, Yunuo Chen, Yin Yang, and Chenfanfu Jiang. 2025. Dress-1-to-3: Single image to simulation-ready 3d outfit with diffusion prior and differentiable physics. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–16.

Weilin Wan, Zhiyang Dou, Taku Komura, Wenping Wang, Dinesh Jayaraman, and Lingjie Liu. 2024. Tlcontrol: Trajectory and language control for human motion synthesis. In *European Conference on Computer Vision*. Springer, 37–54.

Chen Wang, Chuhao Chen, Yiming Huang, Zhiyang Dou, Yuan Liu, Jiatao Gu, and Lingjie Liu. 2025. PhysCtrl: Generative Physics for Controllable and Physics-Grounded Video Generation. In *NeurIPS*.

Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. 2024. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4389–4398.

Haitao Yang, Yuan Dong, Hanwen Jiang, Dejia Xu, Georgios Pavlakos, and Qixing Huang. 2025. Atlas Gaussians Diffusion for 3D Generation. In *The Thirteenth International Conference on Learning Representations*. https://openreview.net/forum?id=H2Gxil855b